Diameter of Cayley graphs generated by transposition trees

Ashwin Ganesan*

Abstract

A problem of much theoretical and practical interest is to determine or estimate the diameter of various families of Cayley graphs. Let Γ be a Cayley graph on $n!$ vertices generated by a transposition tree on vertex set $\{1, 2, \ldots, n\}$. In an oft-cited paper [1] (cf. also [17, p. 188]), it is shown that the diameter of Γ is bounded as

$$\mathrm{diam}(\Gamma) \leq \max_{\pi \in S_n} \left\{ c(\pi) - n + \sum_{i=1}^{n} \mathrm{dist}_T(i, \pi(i)) \right\},$$

where the maximum is over all permutations in S_n, $c(\pi)$ denotes the number of cycles in the disjoint cycle representation of π, and dist_T is the distance function on pairs of vertices of the tree. Observe that evaluating this upper bound requires $\Omega(n!n^2)$ computations since the quantity in braces above needs to be evaluated for each permutation in S_n. In this work, we describe an algorithm to estimate the diameter of the Cayley graph which requires far fewer computations, and furthermore, we show that the value computed by the proposed algorithm is at least as good as (i.e. is less than or equal to) the above upper bound, and that sometimes the value computed is strictly less than the above upper bound. This result is possible because our algorithm works directly with the transposition tree on n vertices and does not require examining any of the permutations. We also show that the value computed by our algorithm is not necessarily unique. We briefly mention a number of new questions and open problems at the end.

Index terms — Cayley graphs; diameter; permutations; transposition trees; combinatorial algorithms; interconnection networks.

*Department of Mathematics, Amrita School of Engineering, Amrita Vishwa Vidyapeetham, Amritanagar, Coimbatore - 641 112, India. Email: ashwin.ganesan@gmail.com.

1. Introduction

Cayley graphs generated by transposition trees were shown in the oft-cited paper by Akers and Krishnamurthy [1] to have diameter that is sublogarithmic in the number of vertices. At that time, the standard topology in interconnection networks against which to compare new topologies with was the hypercube, which has diameter that is equal to the logarithm of the number of vertices. Hence, such Cayley networks were seen as a superior alternative to hypercubes for the topology of interconnection networks, and since then much further work has been done in understanding the properties of such graphs. It is now known that Cayley networks also possess other desirable features such as optimal fault tolerance [2], algorithmic efficiency [3], optimal gossiping protocols [5], and distributed routing algorithms [9]. For further details we refer the interested reader to Du and Hsu[11], Heydemann [18], Lakshmivarahan et al [22] and Xu [23].

The diameter of a Cayley network represents the maximum communication delay between nodes in a network, and obtaining algorithms or bounds for the diameter of various families of Cayley graphs is a research problem of much theoretical and practical interest. The oft-cited paper by Akers and Krishnamurthy [1] provides an upper bound for the diameter of Cayley graphs that are generated by transposition trees. Evaluating this upper bound requires on the order of $\Omega(n!n^2)$ computations, and our main result is an algorithm that much more efficiently computes an estimate for the diameter of the Cayley graph, while still performing at least as well as the previous upper bound, i.e. we show that the value computed by our algorithm is less than or equal to the previously known upper bound. We also show that the value computed by our algorithm is not necessarily unique, though such examples are quite rare. We also discuss a number of open questions and problems that arise due to the algorithm and results given here.

2. Notation and terminology

Given a group G and a set of generators S for the group, the Cayley graph (or Cayley diagram) of G with respect to S is a graph with vertex set G, and with an arc from g to gs whenever $g \in G, s \in S$ (cf. Bollobás [8]). The diameter of the Cayley graph is thus exactly the diameter of the group - the diameter of a group is defined to be the minimum length of an expression for a group element in terms of the generators, maximized over all group elements. The problem of determining this diameter is difficult and remains open even for simple families of Cayley graphs. For example, the problem of determining the diameter of the Cayley graph generated by cyclically adjacent transpositions was studied in Jerrum [19] and remains open. The pancake flipping problem, which corresponds to determining the diameter of a particular permutation group, was studied in Gates and Papadimitriou [14], and while some bounds were given there and a small improvement was provided recently [10], this problem remains open as well.

Let S be a set of transpositions of $\{1, 2, \ldots, n\}$. Construct the transposition graph

$T = T(S)$, whose vertices are $\{1, 2, \ldots, n\}$ and with two vertices i, j being adjacent in T whenever $(i, j) \in S$. It is well known that set of transpositions generates the entire symmetric group S_n if and only if the transposition graph contains a tree [15]. Throughout this work, we focus on the case where the transposition graph is a tree. With a slight abuse of notation, we use the same symbol (i, j) to refer to both a transposition in the set S as well as an edge of the transposition tree. Construct the Cayley graph Γ generated by this transposition tree; the vertices of this graph are the elements of the permutation group S_n generated by this tree, and two permutations π and τ are adjacent in Γ whenever there is a transposition $s \in S$ such that $\pi = \tau s$. Since S is closed under inverses, $(\pi, \pi s)$ is an arc if and only if $(\pi s, \pi)$ is an arc, and hence we assume the Cayley graph Γ to be undirected. Throughout this work, Γ represents a Cayley graph generated by some transposition tree T. We assume throughout that $n \geq 5$ since the problem is easily resolved for all smaller trees.

We let $\mathrm{dist}_T(i, j)$ denote the distance in T between vertices i and j, and $\mathrm{diam}(T)$ denotes the diameter of the tree T.

3. Prior work, and our contributions

We now mention the previous results from the literature on the diameter of Cayley graphs generated by transposition trees as well as summarize our contributions to this problem.

The previous bound on the diameter of Cayley graphs generated by transposition trees is as follows:

Theorem 1. *[1] Let Γ be the Cayley graph generated by a transposition tree T. Then, for any $\pi \in S_n$,*

$$\mathrm{dist}_\Gamma(I, \pi) \leq c(\pi) - n + \sum_{i=1}^{n} \mathrm{dist}_T(i, \pi(i)),$$

where $c(\pi)$ is the number of cycles (including fixed points) in the disjoint cycle representation of π.

Since Γ is vertex-transitive and $\mathrm{dist}_\Gamma(\pi, \tau) = \mathrm{dist}_\Gamma(I, \pi^{-1}\tau)$, by taking the maximum over both sides of the above inequality, we obtain:

Corollary 2. *[17, p.188]*

$$\mathrm{diam}(\Gamma) \leq \max_{\pi \in S_n} \left\{ c(\pi) - n + \sum_{i=1}^{n} \mathrm{dist}_T(i, \pi(i)) \right\} =: f(T).$$

In the sequel we shall often refer to $f(T)$ as the the previously known upper bound on the diameter of the Cayley graph.

We now recall from [1] some terminology that we will be convenient to use. Suppose we are given a transposition tree T on vertex set $\{1, 2, \ldots, n\}$ and a permutation $\pi \in S_n$ for which we wish to determine $\mathrm{dist}_\Gamma(I, \pi)$. At each vertex i of the tree, we

place a marker labeled $\pi(i)$. Thus, the permutation π represents the current position of the markers $1, 2, \ldots, n$ on the tree. To apply an edge (i, j) of the tree to the current position of markers is to say that we switch the markers at the endpoints of the edge (i, j). Note that the permutation corresponding to the new position of the markers is exactly $\pi(i, j)$ (here, we read products or compositions of permutations from right to left). The Cayley graph Γ represents the state transition diagram of the current position of markers, and the problem of determining $\text{dist}_\Gamma(I, \pi)$ is thus equivalent to that of determining the minimum number of edges necessary to 'home' each marker i to vertex i of the tree. This problem is sometimes also referred to 'sorting' a permutation on a tree. Such problems also arise in routing problems in Cayley networks; for example, a node σ receiving a message destined to node τ of Γ, or equivalently, a node $\pi = \tau^{-1}\sigma$ receiving a message destined to node I, needs to figure out which of its neighbors in Γ is closest to the destination node, and this amounts to determining which edge of the tree is optimal in terms of the objective of sorting the current permutation of markers using the minimum number of edges.

The exact diameter value of Cayley graphs generated by transposition trees is known in only some special cases. For example, if the transposition tree is a path graph on n vertices, the corresponding Cayley graph is called a bubble-sort graph. It is well known that the diameter of this Cayley graph is equal to the maximum number of inversions of a permutation, which is $\binom{n}{2}$ (cf. [1], [4], [20]). When the transposition tree is a star $K_{1,n-1}$, the Cayley graph is called a star graph, and it has diameter equal to $\lfloor 3(n-1)/2 \rfloor$ (cf. [1]). For the general case of arbitrary trees, only bounds such as Corollary 2 are known.

Our main results are as follows. We provide an algorithm (Algorithm A below) which more efficiently computes, for any given transposition tree, an estimate of the diameter of the Cayley graph generated by the tree. Our algorithm requires far fewer computations than does the evaluating the previous upper bound given in Corollary 2, and furthermore, we show that the value computed by our algorithm is at least as good as (i.e. is less than or equal to) the previous upper bound given in Corollary 2. We show that sometimes the value obtained by our algorithm is better than (i.e. is strictly less than) the previous upper bound (Theorem 7). Though the value computed by our algorithm is often exactly equal to the previous upper bound, we illustrate some advantages of our algorithm over the previous upper bound; for example, the proofs related to the worst case performance (of $n-4$) of our algorithm are much simpler than those of the previous upper bound (cf. Proposition 8 and the remarks preceding it). It is important to note that we prove that the value obtained by our algorithm is not necessarily unique (Theorem 6). We then discuss some of the open and interesting questions that arise from the proposed algorithm, results, and examples.

Despite the fact that the paper [1] containing the previous upper bound is oft-cited in the literature (this 1989 paper was widely extended in the community immediately), to the best of our knowledge we believe our results here are new.

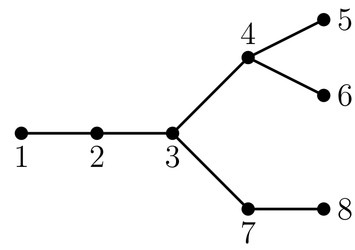

Figure 1: The transposition tree on 8 vertices in Example 1.

4. Algorithm

In this section we describe a combinatorial algorithm that takes as its input a transposition tree T on n vertices and provides as output an estimate of the diameter of the Cayley graph (on $n!$ vertices) generated by T. In Section 5 we prove that the value obtained by our algorithm is at least as good as (i.e. is less than or equal to) the previously known upper bound on the diameter of the Cayley graph generated by the tree.

The notation used to present our algorithm should be self-explanatory and is similar to that used in Knuth [21].

Algorithm A
Given a transposition tree T, this algorithm computes a value β which is an estimate for the diameter of the Cayley graph generated by T. $|V(T)|$ denotes the current value of the number of vertices in T; initially, $V(T) = \{1, 2, \ldots, n\}$.
A1. [Initialize.]
Set $\beta \leftarrow 0$.
A2. [Find two vertices i, j of T that are a maximum distance apart.]
Find any two vertices i, j of T such that $\text{dist}_T(i, j) = \text{diam}(T)$.
A3. [Update β, and remove i, j from T.]
Set $\beta \leftarrow \beta + (2\,\text{diam}(T) - 1)$, and set $T \leftarrow T - \{i, j\}$. If T still has 3 or more vertices, return to step A2; otherwise, set $\beta \leftarrow \beta + |V(T)| - 1$ and terminate this algorithm. ∎

Example 1. Consider the transposition tree $\{(1, 2), (2, 3), (3, 4), (4, 5), (4, 6), (3, 7), (7, 8)\}$ shown in Figure 1. If Algorithm A picks the sequence of vertex pairs during step A2 to be $\{1, 8\}, \{5, 7\}$ and $\{2, 6\}$, then the value returned by the algorithm is $\beta = 7 + 5 + 5 + 1 = 18$. On the other hand, if Algorithm A picks the vertex pairs to be $\{1, 5\}, \{6, 8\}$ and $\{2, 7\}$, then the value returned by the algorithm is still $\beta = 7 + 7 + 3 + 1 = 18$. In this example, the value returned by the algorithm is unique even though the subtrees $T - \{1, 8\}$ and $T - \{1, 5\}$ are non-isomorphic and even have different diameters. ∎

Despite the outcome in the above example where the final value computed by Algorithm A is unique, we prove later that there do exist trees for which the final value computed by the algorithm is not unique (i.e. the final value depends on which

vertex pairs were chosen during step A2), though this non-uniqueness property is rare.

The center of a graph is defined to be the set of vertices of minimum eccentricity, where the eccentricity of a vertex u is defined to be the maximum value of the distance from u to a vertex of the graph. It is well known that the center of a tree is either a single vertex or two adjacent vertices. Also, every path of maximum length in a tree passes through its center.

One way to implement step A2, which picks any two vertices of the tree that are a maximum distance apart, is as follows. Start with an arbitrary vertex u of the tree, and do a depth-first search to find a vertex i farthest that is from u (i and u will be on different 'sides' of the center). Then start at vertex i and do another depth-first search to find a vertex j that is farthest from i. Then, the resulting $i - j$ path has maximum length in the tree.

5. Upper and lower bounds

We first show that the value obtained by Algorithm A is less than or equal to the previously known upper bound on the diameter of the Cayley graph.

Theorem 3. *Let T be a transposition tree on vertex set $\{1, 2, \ldots, n\}$, and let β be the value obtained by Algorithm A for this tree. Then, β is less than or equal to the previously known upper bound on the diameter of the Cayley graph, i.e.*

$$\beta \leq \max_{\pi \in S_n} \left\{ c(\pi) - n + \sum_{i=1}^{n} \mathrm{dist}_T(i, \pi(i)) \right\}.$$

Proof. Let $\{i_1, j_1\}, \{i_2, j_2\}, \ldots, \{i_r, j_r\}$ be the vertex pairs chosen by Algorithm A during the r iterations of step A2, where $r = \lfloor (n-1)/2 \rfloor$. We now construct a permutation π as follows. If n is odd, then T contains only one vertex, i_{r+1} say, when the algorithm terminates. In this case, we let $\pi = (i_1, j_1) \ldots (i_r, j_r)(i_{r+1}) \in S_n$. If n is even, then T contains two vertices, i_{r+1} and j_{r+1} say, when the algorithm terminates. In this case, we let $\pi = (i_1, j_1) \ldots (i_{r+1}, j_{r+1}) \in S_n$. In either case, $r + 1 = \lceil n/2 \rceil$, the value of β computed by the algorithm equals

$$\beta = \left(\sum_{\ell=1}^{r} \{2\, \mathrm{dist}_T(i_\ell, j_\ell) - 1\} \right) + \{(n+1) \mod 2\},$$

and the quantity $f_T(\pi) := c(\pi) - n + \sum_{i=1}^{n} \mathrm{dist}_T(i, \pi(i))$ evaluates to

$$f_T(\pi) = (r+1) - n + \left(2 \sum_{\ell=1}^{r} \mathrm{dist}_T(i_\ell, j_\ell) \right) + 2\{(n+1) \mod 2\}.$$

A quick check shows that the two expressions for β and $f_T(\pi)$ are equal. Hence, for every sequence of vertex pairs chosen by Algorithm A, there exists a permutation π such that the value β returned by Algorithm A is at most $f_T(\pi)$. Hence, $\beta \leq \max_{\pi \in S_n} f_T(\pi)$. ∎

We have not established that the value computed by Algorithm A is unique (and in fact, it isn't sometimes); for a given tree, there can exist more than one pair of vertices that are a maximum distance apart, and different vertex pairs chosen during step A2 can sometimes yield different output values for Algorithm A. Since each of the possible values computed by Algorithm A is at most the previous upper bound, it follows immediately that the largest of the possible values computed by Algorithm A, denoted by $\beta_{\max}$, is also at most the previous upper bound. We now show that $\beta_{\max}$ is also an upper bound on the true diameter value of the Cayley graph Γ:

Theorem 4. *Let Γ be the Cayley graph generated by a transposition tree T. Let $\beta_{\max}$ denote the maximum possible value returned by Algorithm A for this tree. Then,*

$$\operatorname{diam}(\Gamma) \;\leq\; \beta_{\max} \;\leq\; \max_{\pi \in S_n} \left\{ c(\pi) - n + \sum_{i=1}^{n} \operatorname{dist}_T(i, \pi(i)) \right\}.$$

Proof. The second inequality has already been proved. We now prove the first inequality. Let $\pi \in S_n$. Suppose that initially each vertex k of the tree has marker $\pi(k)$. It suffices to show that all markers can be homed to their respective vertices using at most $\beta_{\max}$ edges of the tree.

Consider the following procedure. Pick any two vertices i, j of T that are a maximum distance apart. We consider two cases, depending on the distance in T between vertex i and the current location $\pi^{-1}(i)$ of the marker i:

Case 1: Suppose that the distance in T between vertices i and $\pi^{-1}(i)$ is at most $\operatorname{diam}(T)-1$. Then marker i can be homed using at most $\operatorname{diam}(T)-1$ transpositions. And then, marker j can be homed using at most $\operatorname{diam}(T)$ edges. Hence, markers i and j can both be homed to leaf vertices i and j, respectively, using at most $2\operatorname{diam}(T)-1$ edges. We now let $i_1 = i$ and $j_1 = j$.

Case 2: Now consider the case where the distance in T between vertices i and $\pi^{-1}(i)$ is equal to $\operatorname{diam}(T)$. Let x be the unique vertex of the tree adjacent to $\pi^{-1}(i)$. In the first sequence of steps, marker $\pi^{-1}(i)$ can be homed to vertex $\pi^{-1}(i)$ using at most $\operatorname{diam}(T)$ edges. The last of these transpositions will home marker $\pi^{-1}(i)$ and place marker i at vertex x, whose distance to i is exactly $\operatorname{diam}(T)-1$. In the second sequence of steps, marker i can be homed to vertex i using at most $\operatorname{diam}(T)-1$ edges. Hence, using these two sequences of steps, markers i and $\pi^{-1}(i)$ can both be homed using at most $2\operatorname{diam}(T)-1$ edges. We now let $i_1 = i$ and $j_1 = \pi^{-1}(i)$.

We now remove from T the vertices i_1 and j_1, and repeat this procedure on $T - \{i_1, j_1\}$ to get another pair $\{i_2, j_2\}$. Continuing in this manner until T contains at most two vertices, we see that all markers can be homed using at most

$$\left\{ \sum_{\ell=1}^{r} \left(2\operatorname{dist}_T(i_\ell, j_\ell) - 1\right) \right\} + \{(n+1) \mod 2\}$$

edges. This quantity is equal to the value β returned by the Algorithm when it chooses $\{i_1, j_1\}, \ldots, \{i_r, j_r\}$ as its vertex pairs during each iteration of step A2, and hence this quantity is at most $\beta_{\max}$. Thus, $\operatorname{dist}_\Gamma(I, \pi) \leq \beta_{\max}$ for all $\pi \in S_n$. ∎

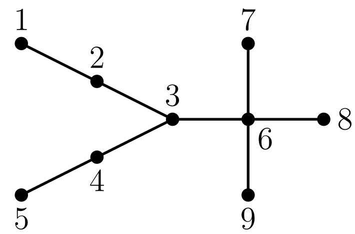

Figure 2: The transposition tree used in the proof of Theorem 6 and Theorem 7.

Let $\mathcal{B}$ denote the set of possible values that can be the output of Algorithm A, and let $\beta_{\max} := \max_{\beta \in \mathcal{B}} \beta$. We have shown that the maximum possible value returned by the algorithm, denoted by $\beta_{\max}$, is an upper bound on the diameter of the Cayley graph. An open problem is to determine whether *each* of the possible values returned by the algorithm is an upper bound on the diameter, i.e. whether β is an upper bound on the diameter of the Cayley graph for each $\beta \in \mathcal{B}$. Our examples so far show that for many families of trees (in fact, for all the ones investigated so far), the minimum possible value returned by the algorithm is also an upper bound on the diameter. We conjecture that the answer to the following problem is in the affirmative:

Problem 5. *Determine whether each of the possible values computed by Algorithm A is an upper bound on the diameter of the Cayley graph.*

6. Non-uniqueness and strictly better performance of Algorithm A

We now show that the value computed by the algorithm is not necessarily unique. Trees such as those given in the non-uniqueness proof below seem quite rare.

Theorem 6. *There exist transposition trees for which the value returned by Algorithm A is not unique.*

Proof. Consider the transposition tree $T_2 = \{(1,2),(2,3),\ (3,6),(3,4),(4,5),\ (6,7),\ (6,8),(6,9)\}$ shown in Figure 2. If Algorithm A picks the sequence of vertex pairs during step A2 to be $\{1,5\},\{2,7\},\{4,8\}$ and $\{3,9\}$, then the value returned by the algorithm is $\beta = 7+5+5+3 = 20$. And if Algorithm A picks the vertex pairs to be $\{1,7\},\{5,8\},\{2,9\}$ and $\{4,6\}$, then the value returned by the algorithm is $\beta = 7+7+5+3 = 22$. Hence, the set of values $\mathcal{B}$ contains $\{20,22\}$. ∎

We now show another advantage of Algorithm A over the previous upper bound - that the value computed by Algorithm A can sometimes be *strictly* less than the previous upper bound on the diameter.

Theorem 7. *The value computed by Algorithm A is always less than or equal to the previously known upper bound on the diameter, and there exist transposition trees for which the value computed by Algorithm A is strictly less than the previous upper bound.*

Proof. The first part of the assertion has been proved earlier. For the second part, consider the transposition tree T_2 shown in Figure 2. It can be confirmed with the help of a computer [16] that the true diameter value of the Cayley graph generated by T_2 is 18 and the previous upper bound $f(T)$ on the diameter evaluates to 22. As mentioned above, if Algorithm A picks the sequence of vertex pairs during step A2 to be $\{1,5\},\{2,7\},\{4,8\}$ and $\{3,9\}$, then the value returned by the algorithm is $\beta = 7+5+5+3 = 20$, which is strictly less than the previous upper bound. ∎

In recent work Ganesan [12], it was shown that the difference between the previous upper bound $f(T)$ and the actual diameter of the Cayley graph Γ is at least $n-4$. The proof given there is quite involved and required an examination of several (over 25) subcases.

As another advantage of Algorithm A over the previous upper bound $f(T)$, we show that the value computed by Algorithm A can also have a difference of at least $n-4$ from the actual diameter value but that the proof of this result is much simpler than the corresponding result for the previous upper bound $f(T)$:

Proposition 8. *For every n, there exists a transposition tree on n vertices, such that the difference between the value computed by Algorithm A and the actual diameter value of the Cayley graph is at least $n-4$.*

Proof. Consider the transposition tree $\{(1,2),(2,3),\ldots,(n-3,n-2),(n-2,n-1),(n-2,n)\}$ shown in Figure 3. The diameter of this tree is $n-2$. After Algorithm A picks and removes two vertices from this tree that are a distance $n-2$ apart, we obtain the path graph on $n-2$ vertices. The unique value computed by Algorithm A is thus $\{2(n-2)-1\}+\{2(n-3)-1\}+\{2(n-5)-1\}+\cdots$, which equals $\{2(n-2)-1\}+\binom{n-2}{2} = \binom{n-1}{2}+n-3$.

Now, any permutation can be sorted on this tree using at most $\binom{n-1}{2}+1$ edges. Indeed, marker 1 can be homed to its vertex using at most $n-2$ edges, and this vertex can then be removed from the tree. Marker 2 can then be homed using at most $n-3$ edges, and so on, and marker $n-4$ can be homed using at most 3 edges. At this point, we arrive at a star $K_{1,3}$, and any permutation on this star can be sorted using at most 4 edges since the diameter of the Cayley graph generated by this star is equal to 4. Thus, the diameter of the Cayley graph generated by this tree is at most $(n-2)+(n-3)+\ldots+3+4 = \binom{n-1}{2}+1$.

Hence, for this transposition tree, the difference between the value computed by Algorithm A and the actual diameter value is at least $n-4$. ∎

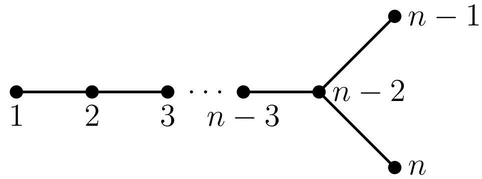

Figure 3: The transposition tree used in the proof of Proposition 8

7. Further questions and open problems

We now discuss some of the new questions and open problems.

Note that our results imply that the value returned by Algorithm A is an upper bound on the diameter for all trees for which $|\mathcal{B}| = 1$ since for such trees any value β computed by the algorithm is equal to $\beta_{\max}$. For such trees, our algorithm efficiently computes a value which is both an upper bound on the diameter as well as better than (or at least as good as) the previously known diameter upper bound:

Corollary 9. *If the transposition tree is such that the value computed by Algorithm A is unique, then this unique value is also an upper bound on the diameter of the Cayley graph.*

An open question is to determine whether this is also the case for the remaining trees. Trees for which $|\mathcal{B}| \geq 2$ are rare, and an open question is to determine whether almost all trees have $|\mathcal{B}| = 1$ ("almost all" in the sense that the proportion of trees on n vertices satisfying this property approaches unity as n approaches infinity):

Problem 10. *Determine whether the value computed by Algorithm A is unique for almost all trees.*

Note that a positive answer to this problem immediately implies that for almost all trees the value computed by Algorithm A is an upper bound on the diameter of the Cayley graph. However, we believe that the following stronger result might also be true (cf. Problem 5):

Conjecture 11. *For each transposition tree, each of the possible values computed by Algorithm A is an upper bound on the diameter of the Cayley graph generated by the tree.*

Even though the value obtained by Algorithm A is sometimes strictly less than the previous upper bound, for all trees examined so far the maximum possible value obtained by Algorithm A is exactly equal to the previous upper bound.

Problem 12. *Determine whether the following statement is true: for all trees T, the maximum possible value $\beta_{\max}$ returned by the algorithm is equal to the previous upper bound $f(T)$.*

One could define the following families of trees. Let $\mathcal{T}$ denote the set of all trees. Let $\mathcal{T}_1 \subseteq \mathcal{T}$ denote the set of all trees for which the sequences of subtrees generated by Algorithm A are isomorphic (i.e. are independent of the choice of vertex pairs

during step A2, unlike the tree shown in Figure 1). Let $\mathcal{T}_2 \supseteq \mathcal{T}_1$ denote the set of all trees for which the sequences of subtrees generated by Algorithm A may or may not be isomorphic, but for which the final value computed by the algorithm is unique (i.e. is independent of the vertex pairs chosen during step A2). Thus, $\mathcal{T}_2$ is the set of trees for which $|\mathcal{B}| = 1$.

Problem 13. *Characterize the families of trees $\mathcal{T}_1$ and $\mathcal{T}_2$.*

8. Concluding remarks

A research area of much theoretical and practical interest is to determine or estimate the diameter of various families of Cayley networks, and this problem remains open even for simple families of Cayley graphs. The fundamental and oft-cited paper [1] provided an upper bound on the diameter of Cayley graphs generated by transposition trees, but evaluating this bound required examining each of the $n!$ permutations of the vertex set of the tree. In this work we described an algorithm to estimate the diameter of the Cayley graph generated by a given transposition tree, and we proved a number of related results. Algorithm A requires far fewer computations to obtain an estimate of the diameter than the previous bound. We showed that the value computed by our algorithm is always less than or equal to the previous upper bound, and that the value computed by our algorithm is sometimes *strictly* less than the previous upper bound. We also showed that the value computed by our algorithm need not be unique, and when it is unique it is an upper bound on the diameter of the Cayley graph. The algorithm, results, and examples given here raise a number of new questions and further problems, some of which have been discussed in the previous sections.

References

[1] S. B. Akers and B. Krishnamurthy. A group-theoretic model for symmetric interconnection networks. *IEEE Transactions on Computers*, 38(4):555–566, 1989.

[2] B. Alspach. Cayley graphs with optimal fault tolerance. *IEEE Transactions on Computers*, 41(10):1337–1339, 1992.

[3] F. Annexstein, M. Baumslag and A. L. Rosenberg. Group action graphs and parallel architectures. *SIAM Journal on Computing*, 19(3):544–569, 1990.

[4] C. Berge. *Principles of Combinatorics.* Academic Press, New York, 1971.

[5] J. C. Bermond, T. Kodate and S. Perennes. Gossiping in Cayley graphs by packets. In *Proceedings of the Franco-Japanese conference Brest July 95*, volume Lecture Notes in Computer Science, 1120, pages 301–315. Springer verlag, 1996.

[6] N. L. Biggs. *Algebraic Graph Theory.* Second Edition, Cambridge University Press, Cambridge, 1994.

[7] N. L. Biggs. *Discrete Mathematics.* Second Edition, Oxford University Press, 2003.

[8] B. Bollobás. *Modern Graph Theory.* Graduate Texts in Mathematics vol. 184, Springer, New York, 1998.

[9] B. Chen, W. Xiao and B. Parhami. Internode distance and optimal routing in a class of alternating group networks. *IEEE Transactions on Computers*, 55(12):1645–1648, 2006.

[10] B. Chitturi, W. Fahle, Z. Meng, L. Morales, C. O. Sheilds, I. H. Sudborough and W. Voit. An (18/11)n upper bound for sorting by prefix reversals. *Theoretical Computer Science*, 410(36): 3372-3390, 2009.

[11] D. Z. Du and F. Hsu. *Combinatorial Network Theory.* Applied Optimization Series, Springer, 1995.

[12] A. Ganesan On the strictness of a bound for the diameter of Cayley graphs generated by transposition trees. *Proceedings of the International Conference on Mathematical Modeling and Scientific Computation*, Springer-Verlag CCIS Series vol. 283, to appear in 2012.

[13] A. Ganesan. Diameter of Cayley graphs of permutation groups generated by transposition trees. `http://arxiv.org/abs/1111.3114`. Technical Report, December 2011.

[14] W. H. Gates and C. H. Papadimitriou. Bounds for sorting by prefix reversal. *Discrete Mathematics*, 27(1):47–57, 1979.

[15] C. Godsil and G. Royle. *Algebraic Graph Theory.* Graduate Texts in Mathematics vol. 207, Springer, New York, 2001.

[16] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.4.12,* `http://www.gap-system.org`. 2008.

[17] G. Hahn and G. Sabidussi (eds.). *Graph Symmetry: Algebraic Methods and Applications.* Kluwer Academic Publishers, Dordrecht, 1997.

[18] M. C. Heydemann. Cayley graphs and interconnection networks. In *Graph Symmetry: Algebraic Methods and Applications*, pages 167–226. Kluwer Academic Publishers, Dordrecht, 1997.

[19] M. Jerrum. The complexity of finding minimum length generator sequences. *Theoretical Computer Science*, 36:265–289, 1985.

[20] D. E. Knuth *The Art of Computer Programming, Volume 3: Sorting and Searching.* Addison-Wesley Publishing Company, 1973.

[21] D. E. Knuth *The Art of Computer Programming, Volume 4A: Combinatorial Algorithms.* Addison-Wesley Professional, 2011.

[22] S. Lakshmivarahan, J-S. Jho and S. K. Dhall. Symmetry in interconnection networks based on Cayley graphs of permutation groups: A survey. *Parallel Computing*, 19:361–407, 1993.

[23] J. Xu *Topological Structure and Analysis of Interconnection Networks*. Kluwer Academic Publishers, 2001.